On Fast Bilateral Filtering using Fourier Kernels

Sanjay Ghosh, *Student Member, IEEE*, and Kunal N. Chaudhury, *Senior Member, IEEE*

arXiv:1603.08081v1 [cs.CV] 26 Mar 2016

***Abstract*—It was demonstrated in earlier work that, by approximating its range kernel using shiftable functions, the non-linear bilateral filter can be computed using a series of fast convolutions. Previous approaches based on shiftable approximation have, however, been restricted to Gaussian range kernels. In this work, we propose a novel approximation that can be applied to *any* range kernel, provided it has a pointwise-convergent Fourier series. More specifically, we propose to approximate the Gaussian range kernel of the bilateral filter using a Fourier basis, where the coefficients of the basis are obtained by solving a series of least-squares problems. The coefficients can be efficiently computed using a recursive form of the QR decomposition. By controlling the cardinality of the Fourier basis, we can obtain a good tradeoff between the run-time and the filtering accuracy. In particular, we are able to guarantee sub-pixel accuracy for the overall filtering, which is not provided by most existing methods for fast bilateral filtering. We present simulation results to demonstrate the speed and accuracy of the proposed algorithm.**

***Index Terms*—bilateral filter, shiftability, Fourier basis, fast algorithm, accuracy.**

I. Introduction

The bilateral filter was introduced by Tomasi and Manduchi in [1] as a non-linear extension of the classical Gaussian filter. The bilateral filter employs a range kernel along with a spatial kernel for performing edge-preserving smoothing of images. Since its introduction, the bilateral filter has found widespread applications in image processing, computer graphics, computer vision, and computational photography [2] - [8].

In this paper, we will consider a general form of the bilateral filter where an arbitrary kernel is used for the range filtering, and a box or Gaussian kernel is used for the spatial filtering [1]. In particular, consider an image $f : I \to \mathbb{R}$, where $I \subset \mathbb{Z}^2$ is a finite rectangular lattice. The output of the bilateral filter $f_{\rm BF} : I \to \mathbb{R}$ is given by

$$f_{\rm BF}(i) = \frac{\sum_{j\in\Omega} w(j)\ \phi\big(f(i-j)-f(i)\big)\ f(i-j)}{\sum_{j\in\Omega} w(j)\ \phi\big(f(i-j)-f(i)\big)}, \quad (1)$$

where $\phi(t)$ is the range kernel and $w(i)$ is the spatial kernel. The spatial kernel is usually a Gaussian [1],

$$w(i) = \exp\left(-\|i\|^2/2\sigma_s^2\right). \quad (2)$$

The window Ω of the spatial kernel is a local neighbourhood of the origin. For example, $\Omega = [-W, W]^2$ for the Gaussian kernel, where $W = 3\sigma_s$. The original proposal in [1] was to use a Gaussian range kernel given by

$$\phi(t) = \exp\left(-t^2/2\sigma_r^2\right). \quad (3)$$

In more recent work, exponential range kernels have been used [9], [10], [11].

Address: Department of Electrical Engineering, Indian Institute of Science, Bangalore, India. Correspondence: kunal@ee.iisc.ernet.in.

The direct computation of (1) requires $O(W^2)$ operations per pixel. In fact, the direct computation is slow for practical settings of W [12]. To address this issue, researchers have come up with various fast algorithms [12] - [17]. While some of these algorithms can reduce the complexity to $O(1)$ operations per pixel for any arbitrary W, there is, however, no available guarantee on the approximation quality that can be achieved using these algorithms. In fact, as reported in [16], a poor approximation can lead to visible distortions in the filtered image. Only recently, a quantitative analysis of Yang's fast algorithm was presented in [18].

In Section II, we recall the idea of constant-time bilateral filtering using Fourier (complex exponential) kernels [19]. In this work, we build on this idea to propose a new algorithm for approximating (1) using the shiftable Fourier basis. The contribution of this work is not the fast algorithm itself, but rather the approximation scheme in Section III, and the subsequent approximation guarantee in Section IV. The approximation scheme can be applied to *any* arbitrary range kernel that has a pointwise-convergent Fourier series. In this respect, we note that all previous approaches based on shiftable approximation were restricted to Gaussian range kernels [16], [19], [20]. We provide some representative results concerning the speed and accuracy of the resulting algorithm in Section V, where we also compare the empirical accuracy of the filtering with the bounds predicted by our analysis.

II. Shiftable Bilateral Filtering

It was demonstrated in [16], [19] that the bilateral filter can be decomposed into a series of Gaussian convolutions using shiftable functions. In particular, since our present interest is in the shiftable complex exponential, consider the function

$$\varphi_N(t) = \sum_{n=-N}^{N} c_n \exp(\iota n\omega t), \quad (4)$$

where $\iota^2 = -1$. By setting (4) as the range kernel $\phi(t)$, we can decompose the numerator in (1) as

$$\sum_{n=-N}^{N} c_n \exp\big(-\iota n w f(i)\big) F_n(i),$$

where

$$F_n(i) = \sum_{j\in\Omega} w(j) f(i-j) \exp\big(\iota n\omega f(i-j)\big). \quad (5)$$

It is clear that a similar decomposition can be obtained for the denominator of (1). We readily recognize (5) to be a Gaussian convolution. As is well-known, the Gaussian convolution in (5) can be efficiently implemented at constant-time complexity (with respect to σ_s) using separability and recursion [21]. In

summary, we can decompose the bilateral filtering into a series of Gaussian filtering. The fast *shiftable* algorithm resulting from this decomposition is summarized in Algorithm 1. We use $G^*(i)$ in line 5 to denote the complex-conjugate of $G(i)$. In line 6, we use $\bar{F}$ and $\bar{G}$ to denote the Gaussian filtering of the images $F(i)$ and $G(i)$. To avoid confusion, we note that the formal structure of Algorithm 1 is somewhat different from that of the shiftable algorithms in [16], [19]. While the cosine and sine components of the complex exponential were used in [16], [19], we work directly with the complex exponential in Algorithm 1. Note that we have abused notation in using $f_{\text{BF}}(i)$ to denote the shiftable approximation of (1) in Algorithm 1.

Data: Image $f : I \to \mathbb{R}$;
Parameters: Filter $w(i)$, and ω, N, and $(c_n)_{-N\le n\le N}$.
Result: Shiftable approximation of (1).
1 **Initialize**: Set $P(i) = 0$ and $Q(i) = 0$ for all $i \in I$;
2 **for** $n = -N, \ldots, N$ **do**
3 $\quad G(i) = \exp(\imath n\omega f(i))$ for $i \in I$;
4 $\quad F(i) = G(i)f(i)$ for $i \in I$;
5 $\quad H(i) = c_n G^*(i)$ for $i \in I$;
6 $\quad$ Compute $\bar{G} = F * w$ and $\bar{G} = G * w$;
7 $\quad P(i) = P(i) + H(i) \cdot \bar{F}(i)$ for $i \in I$;
8 $\quad Q(i) = Q(i) + H(i) \cdot \bar{G}(i)$ for $i \in I$;
9 **end**
10 Set $f_{\text{BF}}(i) = P(i)/Q(i)$ for all $i \in I$.

Algorithm 1: Shiftable Bilateral Filtering.

If the range kernel is not shiftable, one can approximate it using a shiftable function. For example, the shiftable raised-cosines were used in [16] to approximate the Gaussian kernel. Shiftable approximation using polynomials was later presented in [19]. More recently, the classical Fourier basis was used for this purpose in [20]. The above approximations, however, come with the following shortcomings:

- They are customized to work with the Gaussian kernel, and cannot be extended to general range kernels, such as the exponential kernel [10], [11]. Even for the Gaussian kernel, the proposal in [20] requires one to compute the coefficients of the Fourier series. This is computationally intensive (e.g., requires numerical integration, or some analytical properties particular to the kernel), and cannot be done on-the-fly. Indeed, the authors in [20] work with an approximation of the Fourier coefficients, which is only valid for small σ_r.
- Notice that, in most applications of the bilateral filter, the argument t in (3) assumes discrete values. This should be taken into consideration while designing the shiftable approximation. The approximations in [16], [20], however, do not necessarily guarantee that the approximation error at these discrete points are within some user-defined tolerance. This makes it difficult to quantify the overall filtering accuracy. In this paper, we propose a rather simple optimization principle, which has an efficient implementation. This provides us with the desired control on the numerical accuracy of the overall filtering.

III. Progressive Fourier Approximation

We now explain how the above shortcomings can be fixed. As noted above, the argument t in (3) takes on the values $|f(i-j)-f(j)|$ as i and j varies over the image. In particular, t takes values in $\Lambda_T = \{0, 1, \ldots, T\}$, where

$$T = \max_{i\in I} \max_{j\in\Omega} |f(i-j) - f(j)|.$$

Thus, T is the dynamic range of the image measured over the window Ω, which is typically smaller than the full dynamic range. We can compute T using the fast algorithm in [22]; the run-time of the algorithm does not depend on the size of Ω. Without loss of generality, we assume that the range kernel $\phi(t)$ is symmetric. The problem is that of approximating $\phi(t)$ using a shiftable function over the half-interval $[0, T]$. We propose to use the shiftable Fourier basis for this purpose. In particular, we fix some order $N \ge 1$, and consider the shiftable function

$$\varphi_N(t) = d_0 + \sum_{n=1}^{N} d_n \cos(n\omega t), \tag{6}$$

where $\omega = \pi/T$. As is well-known, using the identity $\cos\theta = (\exp(\imath\theta)+\exp(-\imath\theta))/2$, we can write (6) as in (4), where $c_0 = d_0$, and $c_{\pm n} = (1/2)d_n$ for $n = 1, \ldots, N$. The key difference with [20] is with respect to the rule used to set the coefficients $d_0, \ldots, d_N$ in (6). These are set to be the standard Fourier coefficients of $\phi(t)$ in [20]. In keeping with the arguments presented in earlier, we take a different approach and instead try to minimize the error $\phi(t) - \varphi_N(t)$ at the discrete points $t \in \Lambda_T$. In particular, we consider the problem of finding $d_0, \ldots, d_N$ that minimizes the gross error

$$\sum_{t\in\Lambda_T} \big(\phi(t) - \varphi_N(t)\big)^2. \tag{7}$$

This is the classical linear least-squares problem, where the unknowns are $d_0, \ldots, d_N$. Indeed, using matrix-notation, we can write (7) as $\|\mathbf{b} - \mathbf{A}\mathbf{d}\|^2$, where $\mathbf{d} = (d_0, \ldots, d_N)$, $\mathbf{b}$ is the discretization of $\phi(t)$ at the points $t \in \Lambda_T$, and the columns of $\mathbf{A}$ are the corresponding discretization of the basis functions in (6). In particular, let us denote

$$\mathcal{E}_N = \min_{\mathbf{d}\in\mathbb{R}^{N+1}} \|\mathbf{b} - \mathbf{A}\mathbf{d}\|^2. \tag{8}$$

The following fact is the basis of our approximation algorithm to be discussed next.

Proposition 3.1 (Decay of Error): Assume that the Fourier series of the range kernel converges pointwise on the interval $[-T, T]$. That is, for $t \in [-T, T]$,

$$\lim_{N\to\infty} \varphi_N(t) = \phi(t),$$

where $d_0, \ldots, d_N$ in (6) are the Fourier coefficients of $\phi(t)$. Then $\mathcal{E}_N$ decays to zero as $N \to \infty$.

Proof: Indeed, let e_N be the error in (7) when $\varphi_N(t)$ is taken to be the N-th order Fourier approximation of $\phi(t)$. Then, by optimality, we have $\mathcal{E}_N \le e_N$. Since, by assumption, $e_N \to 0$ as $N \to \infty$, the proposition follows. ∎

We note that the Fourier series converges pointwise for any continuously-differentiable function, e.g., Gaussian and polynomials. Convergence is also guaranteed for functions that

are continuous and piecewise-differentiable [23], such as the exponential. Thus, the assumption in Proposition 3.1 covers the commonly used kernels [1], [10], [11].

Proposition 3.1 suggests the following numerical scheme: We fix some user-defined tolerance ε^2. We begin with $N = 1$, and solve (8) to get $\mathcal{E}_N$. If $\mathcal{E}_N < \varepsilon^2$, we stop. Else, we increase N by one and proceed, until $\mathcal{E}_N \leq \varepsilon^2$. In other words, we solve a series of least-squares problems, where the basis matrix $\mathbf{A}$ at each step is obtained by augmenting the $\mathbf{A}$ in the previous step. The whole process can be efficiently implemented using a recursive version of the modified QR algorithm [24]. The main idea is that (8) can be computed by solving $\mathbf{Rd} = \mathbf{Q}^T\mathbf{b}$ using back-substitution, where $\mathbf{A} = \mathbf{QR}$ is the QR-decomposition of $\mathbf{A}$. In the recursive computation, $\mathbf{Q}, \mathbf{R}$, and $\mathbf{Q}^T\mathbf{b}$ at each iteration is computed from the corresponding quantities in the previous iteration using cheap operations. An adaptation of this recursive algorithm to our problem is provided in Algorithm 2. In steps 3 and 14, we discretize the kernel and the incoming Fourier basis. In step 18, r_k denotes the k-th component of $\mathbf{r}$.

IV. Filtering Accuracy

Suppose we are given a range kernel $\phi(t)$ and tolerance ε. We compute the approximation order N and the corresponding coefficients $d_0, \ldots, d_N$ using Algorithm 2. This gives us the corresponding kernel $\varphi_N(t)$ in (4), which is used to approximate (1) using Algorithm 1. In particular, the approximation provided by Algorithm 1 is given by

$$\hat{f}_{\mathrm{BF}}(i) = \frac{\sum_{j\in\Omega} w(j)\varphi_N\big(f(i-j) - f(i)\big)f(i-j)}{\sum_{j\in\Omega} w(j)\varphi_N\big(f(i-j) - f(i)\big)}. \quad (9)$$

By construction, for all $t \in \Lambda_T$,

$$|\phi(t) - \varphi_N(t)| \leq \varepsilon. \quad (10)$$

Similar to [18], we consider the ℓ_∞ (worst-case) error

$$\|f_{\mathrm{BF}} - \hat{f}_{\mathrm{BF}}\|_\infty = \max\big\{|f_{\mathrm{BF}}(i) - \hat{f}_{\mathrm{BF}}(i)| : \ i \in I\big\}. \quad (11)$$

Our goal is to bound (11), which provides us with an estimate of the pixelwise difference between the outputs of the exact and the approximate bilateral filter. In fact, a simple analysis (cf. Appendix) give us the following result.

Proposition 4.1 (Filtering Accuracy):

$$\|f_{\mathrm{BF}} - \hat{f}_{\mathrm{BF}}\|_\infty \leq \frac{2T\varepsilon}{w(0) - \varepsilon}. \quad (12)$$

In other words, the filtering error is essentially within a certain factor of the kernel approximation error ε. To arrive at (12), we have assumed that the weights of the spatial filter add up to unity. Indeed, this assumption can be made since the spatial filter appears in both the numerator and denominator of (1) and (9).

V. Simulation and Conclusion

All simulations reported here were performed using Matlab 8.4 on a MacBook Air with 1.3 GHz Intel Core i5 processor and 4 GB memory. The typical run-time of Algorithm 2 was between 1-15 milliseconds (depending on the order N) for the simulations reported in this section. This is a small fraction of the overall run-time of Algorithm 1. Indeed, the time required to filter a single 512×512 image with a Gaussian kernel is already about 20 milliseconds. In Figure 1, we give an example of the approximation result obtained using Algorithm 2 with $\varepsilon =$ 1e-3. In Figure 2, we compare the coefficients obtained using Algorithm 2 with that obtained by expanding the raised-cosines [16] into (4). Notice that the former decays much more rapidly and hence requires fewer terms.

Data: Kernel $\phi(t)$, half-period T, and tolerance ε.
Result: N and $\mathbf{d} \in \mathbb{R}^N$.

1 $\omega = \pi/T$;
2 $\Lambda_T = \{0, 1, \ldots, T\}$;
3 $\mathbf{b} = [\phi(t)]_{t\in\Lambda_T} \in \mathbb{R}^{T+1}$;
4 **Initialize**: $N = 1$;
5 Set $\mathbf{a} \in \mathbb{R}^{T+1}$ to be the all-ones vector;
6 $\mathbf{A} = \mathbf{a}$;
7 $\mathbf{R} = \|\mathbf{a}\|$;
8 $\mathbf{Q} = \mathbf{a}/\mathbf{R}$;
9 $\mathbf{p} = \mathbf{Q}^T\mathbf{b}$;
10 $\mathcal{E} = \|\mathbf{b} - \mathbf{Qp}\|$;
11 **while** $\mathcal{E} > \varepsilon$ **do**
12 Set $\mathbf{r} \in \mathbb{R}^N$ to be the all-zeros vector;
13 $N = N + 1$;
14 $\mathbf{a} = [\cos(N\omega t)]_{t\in\Lambda_T} \in \mathbb{R}^{T+1}$;
15 $\mathbf{A} = [\mathbf{A} \mid \mathbf{a}]$;
16 **for** $k = 1, \ldots, N-1$ **do**
17 Set $\mathbf{q}_k$ to be the k-th column of $\mathbf{Q}$;
18 $r_k = \mathbf{a}^T\mathbf{q}_k$;
19 $\mathbf{a} = \mathbf{a} - \mathbf{r}_k\mathbf{q}_k$;
20 **end**
21 $r_N = \|\mathbf{a}\|$;
22 $\mathbf{a} = (1/r_N)\mathbf{a}$;
23 $\mathbf{Q} = [\mathbf{Q} \mid \mathbf{a}] \in \mathbb{R}^{(T+1)\times N}$;
24 $\mathbf{p} = [\mathbf{p} \mid \mathbf{a}^T\mathbf{b}] \in \mathbb{R}^N$;
25 Add a row of zeros to $\mathbf{R}$;
26 $\mathbf{R} = [\mathbf{R} \mid \mathbf{r}]$;
27 Solve $\mathbf{Rd} = \mathbf{p}$ using back-substitution;
28 $\mathcal{E} = \|\mathbf{b} - \mathbf{Ad}\|$;
29 **end**

Algorithm 2: Progressive Fourier Approximation.

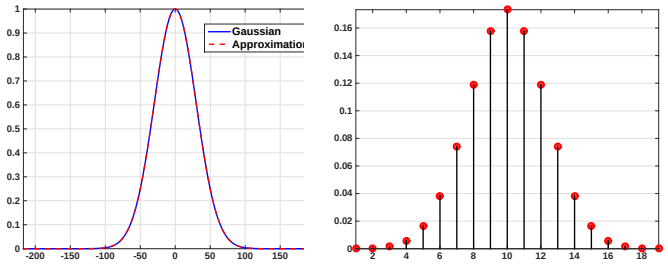

(a) (b)

Fig. 1. Left: Target Gaussian ($\sigma_r = 30$) and the Fourier approximation ($N = 9$) obtained using Algorithm 1. Right: Coefficients $c_{-9}, \ldots, c_9$.

We present some results on the *Barbara* image for which T was computed to be 217. We note that the run-time of

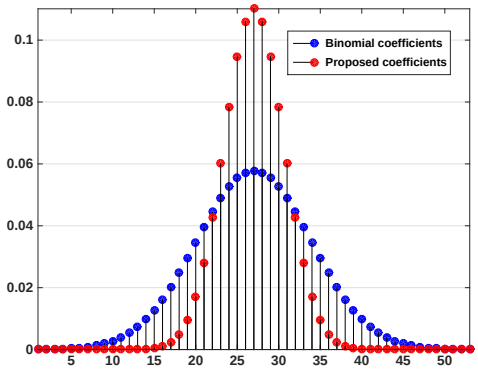

Fig. 2. Comparison of the normalized binomial coefficients from [16] with that obtained using Algorithm 2 ($\varepsilon =$ 1e-3) when $\sigma_r = 10$.

the direct implementation of (1) depends only on the image size and σ_s. On the other hand, the run-time of the proposed algorithm depends on T, tolerance ε, image size, and σ_r. The fact that the run-time is almost independent of σ_s (constant-time algorithm) is evident from the results in Table I. The small fluctuations are essentially due to the variable padding required to handle the boundary conditions for the spatial filtering.

TABLE I
RUN-TIME FOR THE 512×512 *Barbara* IMAGE AT DIFFERENT σ_s AND FIXED $\sigma_r = 30$. THE RUN-TIME OF THE DIRECT IMPLEMENTATION IS 95 SECONDS.

σ_s	1	2	5	8	10	12
Fast ($\varepsilon =$ 1e-3)	630ms	635ms	638ms	640ms	645ms	650ms

The run-time of the proposed algorithm scales inversely with σ_r, which was also observed for the shiftable filtering in [16], [22], [20]. In particular, as σ_r gets small, the Gaussian range kernel tends to a Dirac-like distribution [23]. As is well-known, the Dirac distribution is formally composed of all frequencies. The implication of this fact is that a large N is required to approximate the kernel for small σ_r, and hence the increase in run-time. This is demonstrated with an example in Table II. However, notice that even for small σ_r, the proposed algorithm is much faster than the exact implementation. For $\sigma_r > 20$, the speedup is by a couple of orders.

TABLE II
SAME AS IN TABLE I, EXCEPT THAT σ_r IS VARIED AND $\sigma_s = 3$. THE RUN-TIME OF THE DIRECT IMPLEMENTATION IS 95 SECONDS.

σ_r	10	15	20	30	50	100
Fast ($\varepsilon =$ 1e-3)	2.1s	1.5s	840ms	634ms	450ms	200ms

In Table III, we present the variation of run-time with tolerance ε for a fixed filter setting. It is seen that the order N and hence the run-time changes rather slowly with ε (almost logarithmically in $1/\varepsilon$). We have however not been able to establish this empirical fact, which is deeply tied to the working of Algorithm 2. We next compare the bound in (12) with the actual ℓ_∞ error for the *Barbara* image in Table IV. We note that the error is within the predicted bound. In fact, we are able to predict sub-pixel accuracy when $\varepsilon <$ 1e-5. The bounds are, however, far from being tight. One of the reasons for this is that we have not incorporated any information about the local intensity distribution into our analysis. Derivation of a tighter bound will require a more sophisticated analysis. The present work is a first step in that direction. To best of our knowledge, with the exception of [13], this is the only approximation algorithm that comes with a provable guarantee on the filtering accuracy.

TABLE III
VARIATION OF THE RUN-TIME WITH ε FOR THE *Barbara* IMAGE WHEN $\sigma_s = 3$ AND $\sigma_r = 30$. ALSO SHOWN IS THE ORDER N.

ε	1e-5	1e-4	1e-3	0.01	0.1
N	12	11	10	8	7
Fast ($\varepsilon =$ 1e-3)	910ms	850ms	780ms	715ms	670ms

TABLE IV
COMPARISON OF THE PREDICTED BOUND AND THE ACTUAL ℓ_∞ ERROR FOR THE *Barbara* IMAGE AT $\sigma_s = 3$ AND $\sigma_r = 30$.

ε	1e-8	1e-5	1e-4	1e-3	0.01
N	15	12	11	10	8
Actual (11)	2.7e-8	1.1-4	9e-4	0.01	0.3
Predicted (12)	2.4e-4	0.2	2.5	29.5	561

VI. ACKNOWLEDGEMENT

This work was supported by the Startup Grant awarded by the Indian Institute of Science. The authors would like to thank the anonymous reviewers for their comments and suggestions.

VII. APPENDIX

In this section, we outline the main steps in the derivation of (12). We write (1) as $f_{\rm BF}(i) = P_1(i)/Q_1(i)$, where

$$P_1(i) = \sum_{j\in\Omega} w(j)\phi\big(f(i-j)-f(i)\big)f(i-j),$$

and

$$Q_1(i) = \sum_{j\in\Omega} w(j)\phi\big(f(i-j)-f(i)\big).$$

Similarly, we write (9) as $\hat{f}_{\rm BF}(i) = P_2(i)/Q_2(i)$, where

$$P_2(i) = \sum_{j\in\Omega} w(j)\varphi_N\big(f(i-j)-f(i)\big)f(i-j),$$

and

$$Q_2(i) = \sum_{j\in\Omega} w(j)\varphi_N\big(f(i-j)-f(i)\big).$$

Then $f_{\rm BF}(i) - \hat{f}_{\rm BF}(i)$ can be expressed as

$$\frac{1}{Q_2(i)}\Big[f_{\rm BF}(i)\big(Q_2(i)-Q_1(i)\big) + P_1(i) - P_2(i)\Big]. \qquad (13)$$

From (10), we have $\|Q_1 - Q_2\|_\infty \le \varepsilon$. On the other hand, note that $\|f_{\rm BF}\|_\infty \le T$. This is because $f_{\rm BF}(i)$ is given by the convex combination of $\{f(i-j) : j \in \Omega\}$. Therefore, from (10), we get $\|P_1 - P_2\|_\infty \le T\varepsilon$.

To obtain a lower-bound for $Q_2(i)$ in (13), we note that

$$Q_1(i) = w(0)\varphi(0) + \text{positive terms} \ge w(0),$$

where we have used the non-negativity of the range and spatial kernels. Therefore, using the inverse triangle inequality, we get

$$|Q_2(i)| \ge Q_1(i) - |Q_2(i) - Q_1(i)| \ge w(0) - \varepsilon.$$

By incorporating the above bounds into (13), we get (12).

REFERENCES

[1] C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," *Proc. IEEE International Conference on Computer Vision*, pp. 839-846, 1998.

[2] E. P. Bennett and L. McMillan, "Video enhancement using per-pixel virtual exposures," *ACM Transactions on Graphics*, vol. 24, no. 3, pp. 845-852, *Proc. ACM Siggraph*, 2005.

[3] H. Winnemoller, S. C. Olsen, and B. Gooch, "Real-time video abstraction," *Proc. ACM Siggraph*, pp. 1221-1226, 2006.

[4] J. Xiao, H. Cheng, H. Sawhney, C. Rao, and M. Isnardi, "Bilateral filtering-based optical flow estimation with occlusion detection," *Proc. European Conference on Computer Vision*, pp. 211-224, 2006.

[5] E. P. Bennett and L. McMillan, "Video enhancement using per-pixel virtual exposures," *ACM Transactions on Graphics*, vol. 24, no. 3, pp 845-852, 2005.

[6] K. N. Chaudhury and K. Rithwik, "Image denoising using optimally weighted bilateral filters: A SURE and fast approach," *Proc. IEEE International Conference on Image Processing*, pp. 108-112, 2015.

[7] B. M. Oh, M. Chen, J. Dorsey, and F. Durand, "Image-based modeling and photo editing," *Proc. Annual Conference on Computer Graphics and Interactive Techniques*, pp. 433-442, 2001.

[8] J. Xiao, H. Cheng, H. Sawhney, C. Rao, and M. Isnardi, "Bilateral filtering-based optical flow estimation with occlusion detection," *Proc. European Conference on Computer Vision*, pp. 211-224, Springer, 2006.

[9] B. K. Gunturk, "Fast bilateral filter with arbitrary range and domain kernels," *IEEE Transactions on Image Processing*, vol. 20, no. 9, pp. 2690-2696, 2011.

[10] K. Al-Ismaeil, D. Aouada, B. Ottersten, and B. Mirbach, "Bilateral filter evaluation based on exponential kernels," *International Conference on Pattern Recognition*, pp. 258-261, 2012.

[11] Q. Yang, "Hardware-efficient bilateral filtering for stereo matching," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 36, no. 5, pp.1026-1032, 2014.

[12] F. Durand and J. Dorsey. "Fast bilateral filtering for the display of high-dynamic-range images," *ACM Transactions on Graphics*, vol. 21, no. 3, pp. 257-266, 2002.

[13] Q. Yang, K. H. Tan, and N. Ahuja, "Real-time $O(1)$ bilateral filtering," *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pp. 557-564, 2009.

[14] S. Paris and F. Durand, "A fast approximation of the bilateral filter using a signal processing approach," *Proc. European Conference on Computer Vision*, pp. 568-580, 2006.

[15] F. Porikli, "Constant time $O(1)$ bilateral filtering," *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1-8, 2008.

[16] K. N. Chaudhury, D. Sage, and M. Unser, "Fast $O(1)$ bilateral filtering using trigonometric range kernels," *IEEE Transactions on Image Processing*, vol. 20, no. 12, pp. 3376-3382, 2011.

[17] K. N. Chaudhury, "Fast and accurate bilateral filtering using Gauss-polynomial decomposition," *Proc. IEEE International Conference on Image Processing*, pp. 2005-2009, 2015.

[18] S. An, F. Boussaid, M. Bennamoun, and F. Sohel, "Quantitative error analysis of bilateral filtering," *IEEE Signal Processing Letters*, vol. 22, no. 2, pp. 202-206, 2015.

[19] K. N. Chaudhury, "Constant-time filtering using shiftable kernels," *IEEE Signal Processing Letters*, vol. 18, no. 11, pp. 651 - 654, 2011.

[20] K. Sugimoto and S. I. Kamata, "Compressive bilateral filtering," *IEEE Transactions on Image Processing*, vol. 24, no. 11, pp. 3357-3369, 2015.

[21] R. Deriche, "Recursively implementing the Gaussian and its derivatives, *Research Report*, INRIA-00074778, 1993.

[22] K. N. Chaudhury, "Acceleration of the shiftable algorithm for bilateral filtering and nonlocal means," *IEEE Transactions on Image Processing*, vol. 22, no. 4, pp. 1291-1300, 2013.

[23] L. Grafakos, *Classical Fourier Analysis*, vol. 2, New York: Springer, 2008.

[24] J. Demmel, *Applied Numerical Linear Algebra*, Society for Industrial and Applied Mathematics, 1997.